EXECUTIVE SUMMARY
OPTICAL NETWORK TESTBEDS (ONT) WORKSHOP

1 WORKSHOP OVERVIEW

The Optical Network Testbeds (ONT) Workshop was held at NASA Ames Research Center (ARC) in Mountain View, California, on August 9-11, 2004. The workshop, the seventh in a series of networking focused workshops hosted by the NASA Research and Education Network (NREN), was cosponsored by National LambdaRail (NLR) and Internet2, and was planned and carried out in cooperation with the Federal Large Scale Networking Coordination Group (LSN) and its teams.

The purpose of the ONT Workshop was to provide a forum for LSN agencies and partner organizations to coordinate optical network research programs, testbed architectures and infrastructures, testbed use policies, and applications development, with the goal of identifying key opportunities for cooperation and collaboration.

ONT was an invitation-only workshop, and was attended by over 100 program managers, technologists, and researchers in the area of optical network research and testbed development and use, including participants from LSN agency and partner organizations, as well as from dozens of universities, national laboratories, institutes, network testbed representatives, network providers, and companies.

The report of the workshop is being submitted to LSN agencies for information and for their consideration of the specific recommendations developed during the course of the workshop. The workshop presentations (click on “Agenda”) and the report (click on “Report”) are available on the NREN Web site at www.nren.nasa.gov/workshop7.

The rest of the executive summary follows, from which readers can get a good understanding of the salient observations and discussions from the workshop. In addition, agency and partner organization program and engineering managers are encouraged to read the entire report, which goes into much more depth on all the areas touched on in this executive summary.

2 RECOMMENDATIONS

Below are the summary top-level recommendations from the three breakout groups (BOGs) at the ONT workshop. For more details and next level recommendations, see the remainder of this executive summary and the complete BOG write-ups which appear later in the full ONT Workshop Report.

Note that the top-level recommendations raise the overall main question out of the workshop: Whether LSN should establish as one of its top priorities during the next three years to educate its agency leaders, direct its technical teams, and coordinate with its partner organizations, in order to assure US global ONT competitiveness in the areas of optical networking research, applications, and interdomain connection.

Also note a key observation made during the workshop, that “optical network” is not a defined concept and means different things to different people. No attempt was made at this workshop to clarify or even identify the several different meanings, but the lack of an agreed definition of “optical network” might lead to problems in the future for ONT program managers. LSN may or may not want to tackle this issue, but should certainly at least alert its agencies to this definitional problem.

2.1 From BOG 1, Testbed Research Agendas and Administration/User Policy Issues

ONT Recommendation 1: Optical networking is the next Internet technology wave, and therefore, as a national priority, the US should establish the government, university and industry partnerships that will assure US competitiveness in this “new” new area.

2.2 From BOG 2, User and Application Requirements, Including End-to-end Issues

ONT Recommendation 2: Agencies and partner organizations should work together aggressively to use real applications as drivers for optical internetworking, and to develop a multi-agency connections-type program to enable applications that require optical networks.

2.3 From BOG 3, Testbed Interconnection Frameworks and Issues

ONT Recommendation 3: Optical network testbed programs should adopt optical network testbed interoperability as a primary goal.

ONT Recommendation 4: A testbed working group should be established by LSN/Internet2 engineering and application managers to define an interdomain framework for the exchange of optical network service and transport layer capabilities between networks and with real users.

ONT Recommendation 5: Expand the optical network testbed user community beyond the early adopters.

3 KEYNOTE

The workshop keynote “Optical Networking Outlook: The Private Sector – Then and Now; The Public Sector – Technology Leadership Role” was given by Franz Birkner of FXB Ventures. A well-known venture capitalist in the telecommunications area, Birkner used many graphs and statistics to paint a stark picture of the telecom boom and bust years of the late 1990s and early 2000s. (See his slides on the ONT Web site under “Agenda.”)

A build rate in 2000 of 75,000 fiber miles per day (!), and widespread analyst predictions in 2000 of continuing superheated growth for the optical networking market (predictions of +45% rate of growth for the following four years!), drove massive investments in optical networking technologies by venture capitalists and equipment vendors, totaling over $20B in 2000 alone.

What happened after 2000? The market collapsed. The fourth quarter of 2000 turned out to be the peak, with optical network component vendor revenues of $3B in that quarter. But less than two years later, those revenues had plummeted over 90%. The boom was dead and buried. Birkner then showed several graphs of historical and projected carrier revenues and capital expenditures versus the optical build excess capacity.

The bottom line? Although robust traffic growth continues, there will not be any real carrier capital expenditures recovery to breakeven for the optical networking technology industry until 2006 at the earliest. And this means that equipment vendor research and development (R&D) as well as venture capital will stay away from the area for at least the next two years.

The implication? For the US to maintain global leadership and competitiveness, the public sector must step up to drive the optical “next generation” in the next 3-5 year period. Agencies and universities must continue and grow the significant and substantial activities of US research and engineering (R&E) advanced networking programs, so that US technologists and researchers can stay at the forefront of this vital area. Fortunately, the public sector is indeed stepping up, in the shape of the ONTs that form the focus of this workshop. The powerful collaborations that are being formed have the opportunity to spearhead the next generation of optical networking technology.

Birkner then went on to provide a list of the technology areas that are most promising for the next few years. He ended the keynote with the observation that optical networking is still in its infancy and will ultimately transform telecommunications. The private sector

$20B per year global market in this area will eventually return. In the meantime, the public sector testbeds must provide the means for all US interested parties to work together to maintain global competitiveness in this area during the next 3-5 years.

4 TESTBED PRESENTATIONS

Following the keynote, the remainder of the first day of the workshop featured invited presentations by nine primary ONTs. Each testbed presentation covered an overview of network infrastructure and interconnectivity, as well as the testbed research agenda, driver applications, schedule, and views on existing and potential collaboration and cooperation with other testbeds. (The testbed slides are available on the workshop Web site under "Agenda.") The nine testbeds and their presenters were the following:

Dynamic Resource Allocation via GMPLS Optical Networks (DRAGON)
(Jerry Sobieski, University of Maryland)

Global Information Grid Evaluation Facility (GIG-EF)
(Hank Dardy, Naval Research Laboratory)

Circuit-switched High-speed End-to-End Transport Architecture (CHEETAH)
(Nagi Rao, Oak Ridge National Laboratory)

The OptIPuter and Quartzite LambdaGrid Testbeds
(Larry Smarr, University of California at San Diego)

TransLight
(Tom DeFanti, University of Illinois at Chicago)

Optical Metro Network Initiative Network Testbed: OMNInet
(Joe Mambretti, International Center for Advanced Internet Research, Northwestern University)

National LambdaRail (NLR)
(Tom West, National LambdaRail)

Internet2 Hybrid Optical and Packet Initiative (HOPI)
(Linda Winkler, Argonne National Laboratory)

DOE UltraScience Net
(Bill Wing, Oak Ridge National Laboratory)

In addition, two "**Quick Hits**" were presented:

Pacific Wave
(David Richardson, University of Washington)

WAN in Lab
(Steven Low, Caltech)

5 PANEL DISCUSSIONS

The second day of the workshop was devoted to three panel discussions on the thematic areas of research agendas, user and applications requirements, and testbed interconnection.

5.1 Panel Session 1: Testbed Research Agendas and Administration/User Policy Issues

Chair, Bill Wing, Oak Ridge National Laboratory
Co-chair, Larry Smarr, University of California at San Diego

The discussions in Panel 1 centered on how to identify and promote the research opportunities arising from the availability of low or no cost fiber and the serendipitous opportunities that could arise when the ONTs begin talking to each other. Some main points:

- The US does not have the clear lead on lambda technology, so international collaboration makes sense. The infrastructure security issues of international collaboration need to be separated from the intellectual exchange of ideas.
- The US does have experience in giving applications their own lambdas (e.g., the Electronic Visualization Laboratory at the University of Illinois Chicago, OptIPuter, OMNInet). This approach is viewed as cost-effective. Roughly, the cost of a single 10 Gigabit Ethernet (Gig-E) lambda is equivalent to 1 Full-Time-Equivalent (FTE) technical staff, 50 Terabytes per year (TB/year) of storage, or 1 Gig-E Internet2 connection, any of which costs less than putting a high performance router or MEMS switch at each end, or building a large computing cluster.
- Now is the time to use the low cost fiber that is available. The entire capital expenditure for National LambdaRail (NLR) is about the same as the cost of a single university laboratory building. The cost of one metro area lambda is about the same as the cost of one grad student. These costs are understandable and should be very attractive to university administrators and Federal program managers.
- Innovate, don't just reinvent. Costs will come down for big pipes. 10-40 Gigabit-per-second (Gbps) lambdas will be here soon. By 2010, we will be streaming TB/s. We don't yet know how to use these pipes! New computers will connect 100 Gbps and higher flows directly off the memory controller to Infiniband. All this will challenge the entire Internet layering architecture which is based on today's technology. Plus we have the opportunity to build in security information assurance from the beginning. This architectural evolution will lead to "the backplane **is** the network."
- Latency, not bandwidth, is the fundamental challenge and opportunity of working directly with lambdas. Latencies in a disk drive and across a lambda are basically identical.
- We should be focusing on providing high-bandwidth end-to-end flow pipes on demand, per lambda, not multiplexed. We need to work on the control plane needed to establish and configure these pipes.
- The Internet control plane is seriously challenged – not robust, and BGP/OSPF routing style is not scalable when considering routing lambdas to support individual applications. The history of the Internet is building up networks by aggregation; this fails miserably when applied to big science and big applications.

- "Lambda on demand"? What does it mean? Networks agile enough to give you a lambda in 30 sec? 1 sec? 30 msec? Researchers need to push this envelope.
- Bottom line? We're still at the bottom of the "S curve" in this area. We can't tell agencies what to do or how to collaborate yet, just that they need to become aware of the opportunities, get involved, bring their real applications, and work together to find the best ways to collaborate. The next two years will be extremely dynamic, with a much greater rate of change, than any we have yet seen in optical internetworking up to now.

5.2 Panel Session 2: User and Application Requirements, Including End-to-End Issues

Chair, Steve Corbato, Internet2/University of Utah
Co-chair, Guy Almes, National Science Foundation

Panel 2 looked at applications that need optical networks and the end-to-end problems they raise. Some main points:

- We need to learn how to scale applications end-to-end across a very broad range of network physical and administrative network sizes:
    - International networks
    - National networks
    - Many, many regional optical networks (RONs)
    - Metro networks
    - Department and laboratory networks
    - End systems and clusters
- We **can** build these testbeds! Dark fiber held by the US R&E community will total over 30K route miles, with about half held by NLR and with the other half held by the RONs.
- What new applications could we have with optical networks? The current median rate of bulk TCP over Abilene is 3 Mbps, and the state-of-the-art is 10-50 Mbps. Transfer rates of 100 times this would allow transfers over a single lambda of 10 TB/day = 4 PB/yr.

5.2.1 Example issues

- Users want to be able to "Program at A, execute at B, use/store/visualize at C."
- Applications want to use APIs to "Open/read/write/close," whether to a local disk or over a global network.
- We want to use Layer 3 addressing to virtualize the organization without losing any of the power of internetworking over interconnected lambdas.
- How to integrate IPv6 in an optically driven local-, wide-, storage-area network (LAN/WAN/SAN) environment?

- Use TCP for control and UDP for data, or a new transport entirely? (e.g., eXplicit Control Protocol-- XCP?)
- Need to stay aware of the security requirements, need AAA (authentication, authorization and accounting) middleware to make all this work in the real world.

5.2.2 Example applications

- **National Science Foundation (NSF): electronic Very Long Baseline Interferometry (eVLBI):**
    - 1-10 Gbps from multiple sources need to be correlated in near real time to "diff" out the white noise, cannot be done without optical network speeds.
- **National Institutes of Health/National Science Foundation (NIH/NSF): Biomedical Informatics Research Network (BIRN):**
    - Want to be able to access the world's most powerful electron microscopes in Japan and Korea in near real time, to correlate with US data and to provide closed loop experiment feedback.
    - For brain surgery, want to be able to provide near real time image of brain model to help guide the surgery.
    - For Alzheimer's research, want to be able to pool MRI data of schizophrenics and severely depressed patients across databases in 15 different institutions.
    - Science portals need to be able to launch science workflows with complete AAA support. Teragrid can accomplish over one year of PC science processing in less than one week.
- **Department of Energy (DOE):**
    - Large Hadron Collider (LHC): Data collection to begin in 2007, each 1 TB transaction every 15 minutes will use up an entire lambda.
    - High Energy Nuclear Physics (HENP) Grids: Many systems issues to support long transactions at optical rates with 1000s of users in 100s of institutes in 10s of countries, will prototype the transition to R&E service-oriented grid global optical network.

5.3 Panel Session 3: Testbed Interconnection Frameworks and Issues

Chair, Walt Kaechele, Department of Defense

Panel 3 discussed the many issues involved in achieving interconnection and interoperability of the existing ONTs. Some main points:

- Optical control plane for grids over multiple optical networks is challenging:
    - Network paths coordinated with CPUs and storage advanced reservations.
    - Deterministic low latency and low jitter.
    - Service requests directly from application to network control plane.
    - Very dynamic use of resources.

    - Near-real-time feedback of network availability and performance.
    - Highly distributed management and control.
    - Generalized Multiprotocol Label Switching (GMPLS): Very solid, lots of interoperable implementations, use it!
- Adopt new ways of doing business at the physical layer:
    - Condo fiber builds: Buy point-to-point fiber (20-year lease), light it yourself.
    - Condo wavelengths: Same idea, except with lambdas.
    - Web services to invoke network services? Do in our own ONTs first!
- Carriers would like to offer dynamically switched optical networks, Gig-E and 10 Gig-E over Synchronous Optical NETwork/Synchronous Digital Hierarchy (SONET/SDH), as well as 10 Gig-E LAN PHY transparent transport. Carriers now have dynamic control planes in their optical networks; currently these control planes are not interoperable, but carriers are working in Optical Internetworking Forum (OIF) and International Telecommunication Union (ITU) on the required optical network interfaces and standards.
- All-optical (OOO) switches are now in use in OptIPuter and in Japan. The promise of all-optical switching is enticing, with cost reductions of 20x, 50x, even 256x perhaps becoming available over the next 10 years. But much prototyping in ONTs is needed.
- What should ONTs and RONs do while waiting for widespread carrier and technology offerings? Proceed with what you can do now! Invite participation of vendors, carriers, technology innovators. Evolve the technical skills of your staffs by working with others in a collaborative environment. Evolve to "morphnets" to get real "early production" applications to participate, not only "early adopters."

5.4 JET roles for optical networks

(Paul Love's report on the JET Roadmap Workshop):

- Focus on defining common services for experimental optical networks (e.g., signaling, monitoring).
- Focus on JET coordination of customer-owned R&E optical networks (i.e., technologies, experimental networks, NLR and FiberCo).

6 BREAKOUT GROUP (BOG) DISCUSSIONS

The workshop concluded with three parallel breakout groups (BOGs) charged with clarifying and recommending key opportunities and relationships for cooperation and collaboration among the ONTs.

6.1 BOG 1: Testbed Research Agendas and Administration/User Policy Issues

Lead: Joe Mambretti, Northwestern University

6.1.1 Observations

There was consensus among BOG 1 attendees that:

- The optical networking research discussed in this BOG represents the next generation of internetworking, following in the footsteps of the government-funded research that led to the Internet and Web.
- A set of special circumstances has created a unique window of opportunity for the US to lead the creation of the next wave of Internet innovation by pursuing advanced optical networking research in Federal and partner ONTs.
- There are three main technology forcing functions at work: (a) orders-of-magnitude **cost-performance** improvements promised by optical network technologies (b) will allow networks to **meet never-before-achievable requirements** of a wide array of large scale science and engineering applications (c) by **enabling terascale system architectures** that will dramatically enhance or in some cases may even replace much of the existing Internet architecture.

6.1.2 Discussions

6.1.2.1 Importance of the Research

- Future R&E networks need to support two kinds of traffic: (a) aggregation traffic (including support for QoS, priority, preemption, etc.), and (b) single-stream traffic (e.g., support for mega instruments, supercomputers, terascale databases, haptics, rapid onset of large streams).
- To do this, future edge networks will need to be able to deliver specialized services and flexible guarantees of performance to individual high performance users and systems located in laboratories and data systems connected directly to the edge. This will require new kinds of user visibility and control into core as well as peer networks.
- Due to the blazing speed of optical networks, making this work will require researchers to substantially simplify network architectures and protocols rather than adding new layers of complexity. This is the work that needs to be done in the Federal and partner ONTs in the next 3-5 years.

6.1.2.2 Importance of the Testbeds

- Optical network testbeds (ONTs) bring optical networking technology cultures together, to interact in real world environments. Technical areas such as optical transport and IP routing/switching are very complex, with great opportunities for synergy.
- Many synergies are not obvious until complex interactions happen among materials, devices, subsystems, systems, integration, operation, troubleshooting, diagnostics and repair. These interactions are what will drive the new understandings and innovation breakthroughs that will come out of the ONTs.
- Not only government and university researchers, but also startups and larger companies, need such “safe haven” testbeds to try out their new ideas in a rigorous but friendly environment. This frequently costs the government very

little beyond establishing the testbed facilities and interconnections, and funding the graduate students who provide the bulk of the skilled low-cost labor.

- Many future US technologists will earn their masters and PhDs and spend their post-docs working on these ONTs and applications. These long-term investments in people are needed to produce the great companies and teachers that will fuel the next generation of US networking competitiveness.

6.1.2.3 Why Not Wait for the Carriers?

- The BOG members revisited and strongly agreed with the points made in the keynote that the US public sector, not the carriers, needs to lead the next wave of internetworking innovation by establishing and strongly promoting cooperation and collaboration among the agency and partner ONTs. For an in-depth understanding of this key rationale issue, see the keynote slides plus sections 2.3 to 2.5 of the full BOG 1 report.
- The US has a 2-4 year window of opportunity to use the more than 30,000 miles of domestic fiber available to Federal and partner R&E networks at low or no cost to establish the prototype implementations of US optical internetworking.
- At Gigabit Ethernet and higher speeds, and with disk-like latencies, scientists will begin to create **entirely new science**. This happened in the 1980s, when science networks changed from 56 kbps dial-up to 10 Mbps Ethernet, and it happened again in the 1990s with the change from 1.5 Mbps T-1 access connections to 155 Mbps OC-3.
- With this two-orders-of-magnitude performance explosion, new bottlenecks will show up in systems and middleware. Research on understanding and then mitigating these bottlenecks is what the ONTs need to focus on in the 3-5 year timeframe. A systems approach to providing real scientists with truly usable lambda networks is therefore key to the 3-5 year ONT research agendas.
- The ONT testbeds are not only the research vehicles for US competitiveness in the next wave of the Internet, but are also the agencies' only affordable growth paths for development of their next-generation production networks.

6.1.2.4 ONT Technology Research Topics

- Important research areas identified in BOG 1 included Layer 3, Layer 4 routing, Layer 2 and 3 switching, wavelength switching, new edge vs. core paradigms, devices and components, management and control planes. See the corresponding sections of the full BOG 1 report for details.

6.1.3 Top-Level BOG 1 Recommendation

Optical networking is the next Internet technology wave, and therefore, as a national priority, the US should establish the government, university and industry partnerships that will assure US competitiveness in this "new" new area. Agencies and partner organizations should establish as one of their top priorities during the next three years to educate their leaders, direct their technical teams, and coordinate with their partner organizations, in order to assure US global ONT leadership in the areas of optical networking research, applications, and interdomain connection.

6.2 BOG 2: User and Application Requirements, Including End-to-End Issues

Lead: Alan Blatecky, Renaissance Computing Institute
Co-lead: Tom DeFanti, University of Illinois at Chicago

6.2.1 Discussions

BOG 2 attendees discussed four broad areas, developing the following consensus outlines that summarize these discussions.

6.2.1.1 Deployment of optical networks for selected research projects

a. What projects or research efforts require optical networks and on what time scales (utilization)? How multidisciplinary? Investigate and evaluate novel technology, tools, diagnostics and security.

b. Where is the cut-off, what is the criteria to justify deployment? One-time and continuing costs should be compared to the fraction of the overall project cost and fees for scaled-up shared networks.

c. Innovation for application and network scaling.

6.2.1.2 Establishment of a multi-agency "optical connections" program

a. Small number of players at outset; ability to scale is important.

b. Driven by applications and need for optical network capabilities.

c. Cross-agency network management and diagnostic tools development and deployment.

6.2.1.3 Development of hybrid network testbeds (optical, wireless and IP)

a. Future networks will not be monolithic.

b. Interoperability and transparency are essential.

c. Network-aware applications/networks that can advise applications.

d. Many applications require such capabilities: sensors, disaster, health, sharing high-performance computing (HPC), outreach, etc.

6.2.1.4 Development of optical "morphnet" networks bridging advanced research and early "production" networks

a. Dual synergistic approach: network research and domain research.

b. Provide entry and access for providers.

c. Applications inform network architects and providers.

d. Technology transfer process more implicit.

During the above discussions, the following areas were covered in more depth:

- Involving the user community.

- Application requirements for resource management and measurement
- Control plane and infrastructure issues
- Other recommendations for Federal agencies.

See the corresponding sections of the full BOG 2 report for a detailed summary of the discussions in these areas.

6.2.2 Top-Level BOG 2 Recommendation

Agencies and partner organizations should work together aggressively to use real applications as drivers for optical internetworking, and to develop a multi-agency connections-type program to enable applications that require optical networks. Optical networks are fast becoming the "new" new thing in high performance internetworking; therefore agencies and partner organizations should provide incentives for real science and application users to get involved in order to effectively drive the optical network testbeds and the research and education (R&E) early production optical networks in this vital area. This will assure that the resulting infrastructure of interconnected networks will meet the true requirements of demanding existing and future "e-science" and government applications.

6.3 BOG 3: Testbed Interconnection Frameworks and Issues

Lead: Linda Winkler, Argonne National Laboratory
Co-lead: Jerry Sobieski, University of Maryland

6.3.1 Observations

There was consensus among BOG 3 attendees that:

- Optical network lightwave-based interconnection services are necessary to meet future terascale applications requirements, because the current router-based Internet architecture is not sufficiently scalable to provisioning of core network support for terascale applications.
- Agency long-term support commitments are required for the advanced network research activities necessary to develop and deploy dynamic, reliable, and robust optical network services needed to attract next-generation science and application users.
- Existing research testbeds are promising, but are missing a key enabling point: interoperability.

6.3.2 Discussions

- Goals of existing and future funded programs should be modified to incorporate a new "interoperability" requirement.
- Optical networks in general, but more specifically the target research and education (R&E) optical network testbeds, don't currently interconnect, much less pass common transport and service layer information.

- There is no common service definition that forms the basis for the exchange of information among interconnected advanced networks.
- To remedy this situation, an intertestbed working group needs to be established to define an interdomain framework for the exchange of service and transport layer capabilities.
- Success of applications that seek to use advanced optical transport and service capabilities requires global reach.
- There does not yet exist a critical mass for commercial viability of these new, advanced optical networking services.
- To achieve this crucial goal, agencies and partner organizations must enter into a sustained (long term) mutual commitment that the new network services provided to users will be dependable, robust, and deterministic.
- Agency basic network research programs need to consider what it would take to give applications terabit-per-second connectivity.

6.3.3 Top-Level BOG 3 Recommendations

1. ONT programs should adopt optical network testbed interoperability as a primary goal. Agencies should adjust (e.g., refine, expand or redefine) existing goals to assure that these interoperations happen. Agencies and partner organizations should work together to develop the multi-agency program relationships needed to interconnect their existing and planned optical network testbeds.

2. A testbed working group should be established by LSN/Internet2 engineering and application managers to define an interdomain framework for the exchange of optical network service and transport layer capabilities between networks and with real users. The working group should identify those domains best suited to begin the task of building this framework for interconnecting independent optical networks and develop some near-term goals to illustrate the benefit of this approach. The framework must enable each domain to describe the types of services it can provide, without requiring domains to provide specific services.

3. Expand the optical network testbed user community beyond the early adopters. To achieve this goal, active involvement of science program funding sources, as well as testbed interconnection with regional and international optical networks and their user communities, are critical. Just as important, there must be a mutual commitment between network providers and the science community that the new network services provided to users will be dependable, robust, and deterministic, so that users will actively embrace the new services and pursue the new science paradigms that will be enabled.